

THE TOP 10 CARS OF 2012

MUSTANG BOSS 302, AUDI A7, BMW 1M...

ENTER
TO

WIN!

\$3,500
IN HOLIDAY GEAR

Science Technology Automotive Home Adventure
PopularMechanics.com December 2011

Popular Mechanics

AMERICA'S NEXT

WAR PLANE

THE X-47B

Radar-Beating,
Carrier-Ready,
Two Tons of Bombs,
& It's a Robot

THE WORLD'S FASTEST COMPUTER

How We're Building It — And
Why It Matters

Hey, check out our iPad app, too!



THE OTHER
CONTENDERS

X-45C

UNMANNED STEALTH
STRIKE AIRCRAFT

F-35C


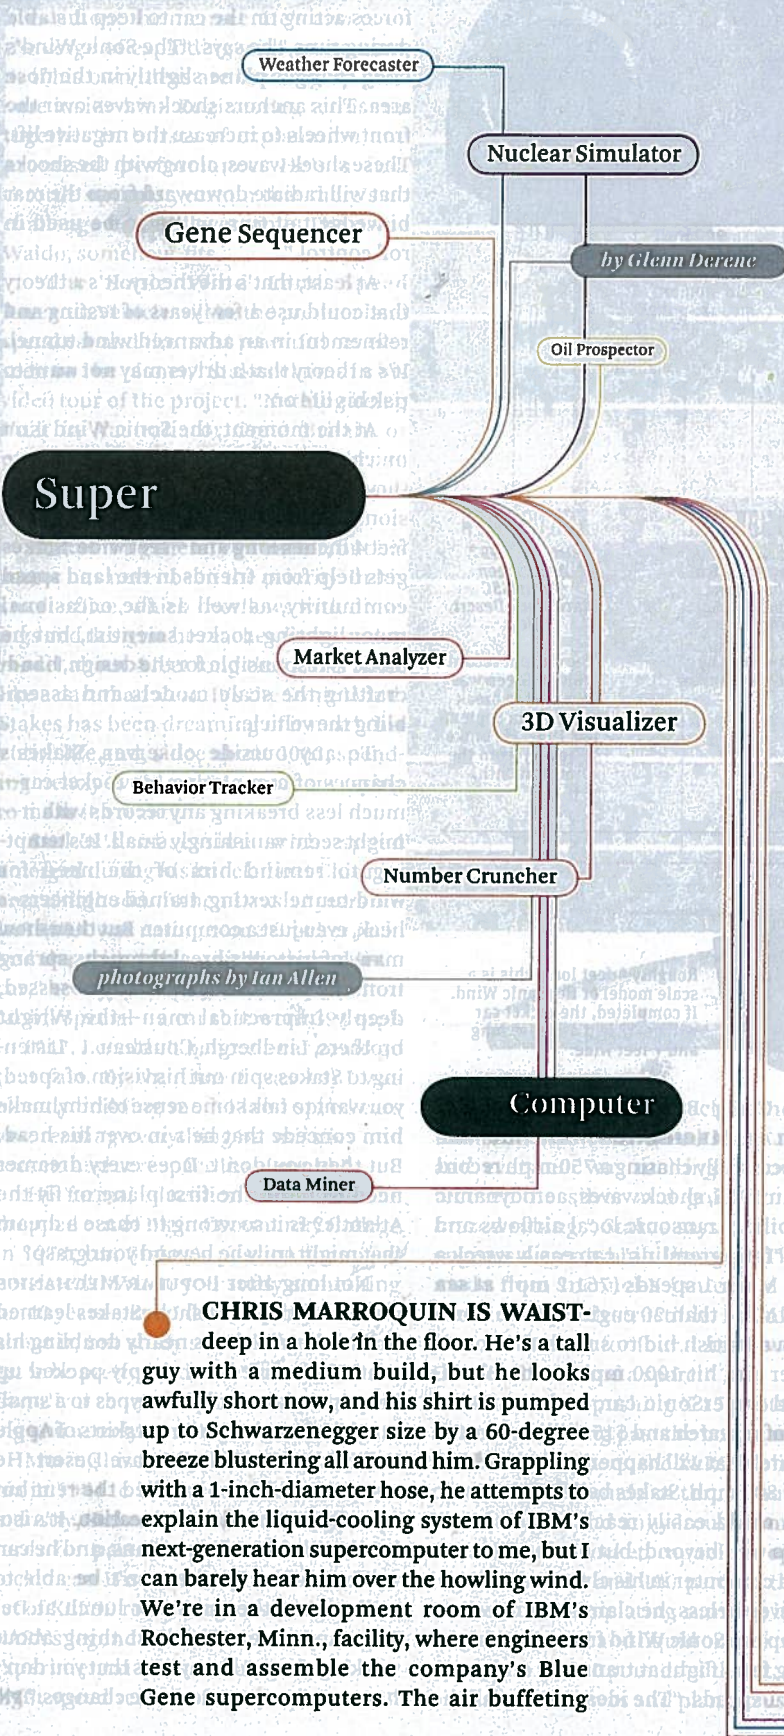
MOST ADVANCED MANNED
FIGHTER JET


F/A-18E/F

FAST, RELIABLE, AND LOADED FOR BEAR

How to Pick the
Safest Winter Tires
12-Volt Drill Test
Shoot Great Smartphone
Movies & More!

diy





Fiberoptic networking and power cables intermingle with flexible rubber hoses carrying cooling water to a prototype Blue Gene/Q supercomputer at IBM's Rochester, Minn., facility.

Anatomy of a Supercomputer

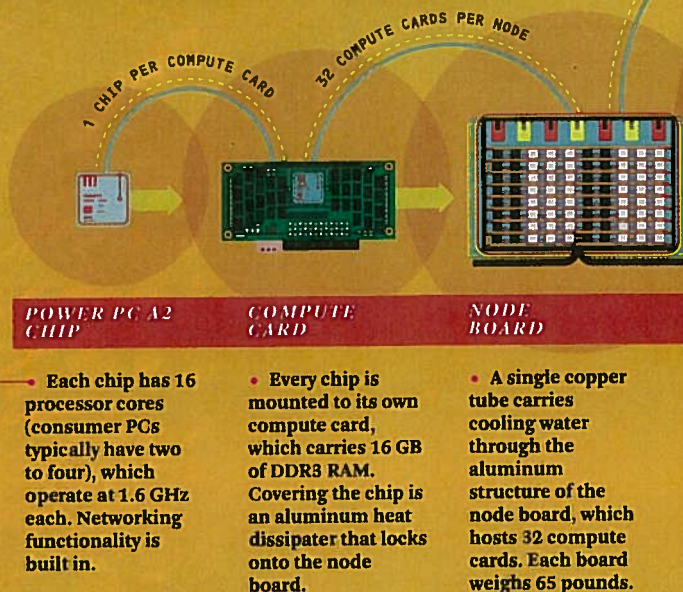
Marroquin cools a small, four-rack Blue Gene/P system capable of 13.9 teraflops per rack, but the hose he's holding is part of a far more advanced cooling system. Filled with deionized water, the anti-corrosive agent benzotriazole and a dose of biocide, the tube feeds into a prototype of the company's new Blue Gene/Q computer. The Blue Gene/Q rack sitting on the raised floor has its own circulatory system—850 feet of copper pipe, with check valves, quick-disconnect rubber hoses and an electronic monitor that measures flow rate, pressure and dew point—designed to shut down if anything goes awry. “You don’t want any drips,” Marroquin says.

As sophisticated as the cooling system is, what launches this machine into the realm of technological superlatives is its processing power: Each rack contains 1024 computer chips, and every one of those chips has 16 processor cores. That’s a total of 16,384 processors, making it capable of 209 teraflops, 15 times more power per rack than the Blue Gene/P. Within the next year IBM will ship 96 Blue Gene/Q racks to Bruce Goodwin at Lawrence Livermore National Laboratory (LLNL) in California. Collectively, those racks will become the most powerful computer in the world. It should be able to predict the path of hurricanes, decode gene sequences and analyze the ocean floor to discover oil. But Goodwin primarily wants to use it to blow up a nuclear bomb.

GOODWIN USED TO EXPLODE NUKES THE old-fashioned way. From 1983 to 1991, he designed and oversaw five nuclear weapons tests at the Department of Energy’s Nevada Test Site. He and other engineers would dig a 2000-foot-deep hole, toss a warhead and some highly specialized monitoring equipment into a 10-story-tall, 1-million-pound iron canister and lower it into the hole. Then everybody would move way the heck back, cross their fingers and detonate. “Sitting in the control room 10 miles away, it felt like a magnitude 5 or 6 earthquake,” Goodwin says.

All that changed in October 1992, when then President George H.W. Bush declared a moratorium on nuclear testing in anticipation of the Comprehensive Nuclear-Test-Ban Treaty of 1996. After that, if the United States wanted to test any of the warheads in its multithousand-weapon arsenal, it had to do a computer simulation. Thus, our interest in really powerful computers was nationalized.


Really powerful computers have been around as long as computers themselves, but the term supercomputer didn’t arrive until 1976, when Seymour Cray built the Cray-1. It cost \$8.8 million (\$35 million in today’s dollars) and cranked up to 160 megaflops. Yesterday’s supercomputer, however, has less power than today’s personal



computer—a modern PC has more than 50 times the processing horsepower of the original Cray. In fact, the “super” prefix is so fuzzy that many computer scientists eschew the term supercomputer altogether and call such machines high-performance computers, or HPCs. In an attempt to bring some clarity to the genre, in 1993 a private group called the Top500 project started publishing a twice-yearly list of the 500 most powerful computers in the world. If your computer is on the list, it is by definition a supercomputer.

For 17 of the Top500 list’s 18 years, the U.S. and Japan have swapped supremacy. But in October 2010, China claimed the top spot with the 2.6-petaflop Tianhe-1A. The computer scientists who design and build these systems tend to work for multinational companies and are cautious about characterizing what they do as a statement of national pride. Regardless, supercomputers have come to symbolize the technological prowess of the countries that build them—a silicon-age version of the space race. In a sign of the whipsaw speed of technological progress, Japan eclipsed China just eight months later, in June 2011, unveiling the 8-petaflop K Computer. The Chinese countered in August, outlining a road map to “exascale” computing, essentially promising a 125-fold increase in computing power within 10 years. If Tianhe-1A was China’s Sputnik moment, exascale is its moonshot.

The supercomputer’s role in maintaining America’s nuclear weapons justifies its status as a national security interest. But China’s challenge to the West’s computing dominance has led many computer scientists and policy wonks to claim that supercomputing is essential to U.S. economic security as well. These machines are force mul-



Top 5 Super-computer Sites

When Sequoia is fully operational in 2012, it is expected to become the most powerful computer on earth. How does it compare?

Location	LIVERMORE, CA	KORE, JAPAN	TIANJIN, CHINA	OAK RIDGE, TN	SHENZHEN, CHINA
Computer	Sequoia	K Computer	Tianhe-1A	Jaguar	Nebulae
Laboratory	Lawrence Livermore National Laboratory	Riken Advanced Institute for Computational Science	National Supercomputing Center	Oak Ridge National Laboratory	National Supercomputing Center

multipliers for American scientists, engineers and businesses, the argument goes, and whoever builds the best ones gains an advantage. Supercomputers don't just reflect intellectual and technological power, they also reinforce it.

THE FOLKS AT IBM ROCHESTER BETRAY little interest in China's goal of supercomputing dominance. Their job is to work out the engineering for Blue Gene/Q, and they deliberately focus on the technology, not the politics. They are classic pocket-protector engineers, and their titles are inelegant bureaucratic artifacts that offer little clue to their actual roles. "We're a very small, roll-up-your-sleeves team effort," says Pat Mulligan, development manager for Global Server Integration (who, for the record, had his sleeves rolled up when we spoke). "We're not overly nationalistic, we just

want to make the best computer we can."

The building where Marroquin, Mulligan and the rest of the IBM team are creating the 21st century's most powerful computers is a monument to mid-20th-century corporate futurism. Designed by architect Eero Saarinen (who also designed the St. Louis Gateway Arch), the sprawling structure is clad in dark blue glass. Hallways a half-mile long stretch through the interior. At some point IBM—always pushing the technological envelope—concealed wires in the hallway floors to guide robots that delivered parts and machinery from one assembly room to another. The robots are long gone, a dream of mechanical efficiency undone by reality: They were slow and

CONTINUED ON PAGE 124

Supercomputer (CONTINUED FROM PAGE 89)

broke down so often that the facility switched to human-guided forklifts.

The Blue Gene/Q computers I'm getting a look at in midsummer are not part of Bruce Goodwin's supercomputer (named Sequoia). These are test models, used to work out the kinks in the hardware and software. The manufacturing of Sequoia's 96 racks was due to ramp up soon after my visit, but Goodwin and his team at Lawrence Livermore are already logging in to Blue Gene/Q and tinkering from afar; a sign on one of the racks in the Rochester assembly room says LLNL REMOTE ACCESS MACHINE.

Goodwin's Terascale Simulation Facility (TSF) at Livermore is one of two DOE centers that perform nuclear simulations as part of the Stockpile Stewardship Program (the other is at Los Alamos National Laboratory in New Mexico). To get a simulation that delivers an acceptable degree of accuracy, Goodwin's team models a 50-microsecond explosion in three dimensions down to a scale of 10 microns. "It gets very complicated," Goodwin says. "These things are imploding and exploding, and you have to track the fluid mechanics with the precision of a Swiss watch." Every time a component is changed or upgraded in a U.S. nuclear warhead, the TSF virtually tests the bomb to make sure it will still go boom. The computer simulations have revealed aspects of nuclear fission that testers hadn't anticipated, and, consequently, the number and complexity of algorithms have increased over time. Modern simulations model only parts of a full explosion, and even then, the most complex sims Goodwin runs use about a million lines of code. If you had 1600 years, the calculations could conceivably be done on a laptop; Livermore's current 500-teraflop Blue Gene/P system, named Dawn, gets a high-complexity sim done in a month. When the 20-petaflop Sequoia system goes live in 2012, the test time should drop to a week.

TO UNDERSTAND SUPER-computers, you need to understand flops, or floating-point operations per second. Flops are essentially math with decimals, as opposed to integer calculations, which require whole numbers. When it comes to hardcore number-crunching, flops are more data-efficient than integers—consider Avogadro's number, expressed as 6.02×10^{23} , compared with its integer alternative, which would fill out most of this sentence. High-performance computers are super-floppers: Sequoia's 20 petaflops equals 20 *quadrillion* calculations per second.

trades that drove the Dow down 600 points in 5 minutes.

The secret to supercomputing is parallel processing. The design of a supercomputer allows the machine to break up a task—say, predicting the path of a tornado—into lots of interdependent calculations, then groups of processors crunch the numbers all at once. To make things even faster, each of Sequoia's chips has onboard networking and can share data directly with any other chip in its rack.

It's a brute-force approach to math, and it is surprisingly powerful. A Blue Gene/P computer recently cal-

The logic for building supercomputers is strikingly similar to that of nuclear deterrence. We must have more computing power than our competitors or they will use their technological superiority against us.

So high-performance computing is predicated on the idea that many of the world's most complicated problems are ultimately reducible to pure math. And those problems range from matters of national security (the viability of Goodwin's nukes) to day-to-day concerns (predicting the weather this weekend—and the weekend after that). Not only are supercomputers routinely used in research (climate modeling, gene sequencing, artificial intelligence), but they are also becoming essential to commercial enterprises such as drug development, oil exploration and aircraft and automotive design, as well as product R&D. For example, Arizona-based Ping has used Cray supercomputers to aid in golf-club design. Supercomputers let companies speed products through the development cycle by virtualizing much of the design and testing. High-performance computing can also have more ominous consequences—Wall Street's "flash crash" in May 2010 was caused by a chain reaction of HPCs making high-frequency

culated pi to the billionth digit. It's also surprisingly scalable. Sequoia will have 96 racks, but Dr. George L.T. Chiu, one of IBM's top HPC scientists, claims that with a few simple hardware and software changes, Blue Gene/Q could theoretically support up to 32,768 racks, with an estimated compute power of 6848 petaflops. "The actual limit is the dollars you're willing to spend," Chiu says. "And, of course, you have to have the power."

Oh, yes, electricity. That's the other big issue with HPCs. Sequoia will be the most powerful supercomputer in the world, but it will also be one of the most power-hungry. At peak load, Sequoia is expected to operate at 9-plus megawatts, enough to power more than 7200 homes. It turns out, however, that Sequoia will also be the world's most power-efficient computer, churning out 2 gigaflops per watt. By comparison, the K Computer in Japan, which operates at 9.9 megawatts, puts out just 800 megaflops per watt—accomplishing only 40 percent of the calculations with the same electricity. But like processing power,

electricity use scales linearly as you add racks to a supercomputer. If you double the racks on Sequoia, you get a computer that's twice as fast—but you also get a computer that's twice as power-hungry. Being the world's most efficient computer helps to mitigate that consumption, but only to a point.

CHINA ISN'T THE ONLY COUNTRY aiming for exascale. The Department of Energy deems it critically important to American technological competitiveness, and companies such as Intel and Nvidia are promising exascale performance by the end of this decade. It's a technological challenge that goes beyond mere improvements in processing power. "To build exascale you have to have a vision of what applications will look like 10, 15 and 20 years from now," says Dave Turek, IBM's vice president of exascale computing. Turek and his contemporaries foresee a future where the volume and speed of data coming at machines like this will be several orders of magnitude higher than it is now, and will require a ground-up re-engineering of some of the fundamentals of computing, such as data storage, networking, software and power systems.

Supercomputing is an expensive hobby for a nation to have. The DOE puts the combined development costs of Sequoia and Blue Gene/P Dawn at about \$250 million. Plus, the annual electric bill to operate a petascale computer runs \$5 million to \$10 million. High-performance-computer scientists know that costs like these can't be allowed to scale along with the gains expected from exaflop

machines. But Goodwin and others in his field see these computers as essential. He points out that China's government has a stated goal of using supercomputers to gain an industrial edge, and we should be doing the same. "We can do all of the engineering 'what ifs' on a supercomputer and bring a product to market five times faster than when you actually had to make things to see if they worked," he says. "Think about what it means to the national economy if Boeing, General Motors or General Electric can get to market in months instead of years. It matters, and if someone can get there five times faster than you, you're going to go out of business."

We justify the expense of these machines today because they help to maintain our nuclear stockpile, but the logic for building them in the future is strikingly similar to that of nuclear deterrence itself. We must have more computing power than our competitors or they will use their technological superiority against us. This made me wonder what kind of computer would be fast enough for Goodwin's 50-microsecond nuclear sims. I asked him: If a 500-teraflop computer could do it in a month, and a 20-petaflop computer could do it in a week, could an exaflop computer do it in real time? "An exaflop machine is way too slow to run such a simulation in real time," he answered via email. He told me a real-time nuclear simulation would require a 100-yottaflop computer—that's 100×10^{24} calculations per second, 100 million times faster than an exaflop machine. Another floating-point operation. **PM**